arXiv:1704.08615v1 [cs.CV] 27 Apr 2017

Saliency Benchmarking: Separating Models, Maps and Metrics

Matthias Kümmerer, Thomas S. A. Wallis, Matthias Bethge

Abstract

The field of fixation prediction is heavily model-driven, with dozens of new models published every year. However, progress in the field can be difficult to judge because models are compared using a variety of inconsistent metrics. As soon as a saliency map is optimized for a certain metric, it is penalized by other metrics. Here we propose a principled approach to solve the benchmarking problem: we separate the notions of saliency models and saliency maps. We define a saliency model to be a probabilistic model of fixation density prediction and, inspired by Bayesian decision theory, a saliency map to be a metric-specific prediction derived from the model density which maximizes the expected performance on that metric. We derive the optimal saliency map for the most commonly used saliency metrics (AUC, sAUC, NSS, CC, SIM, KL-Div) and show that they can be computed analytically or approximated with high precision using the model density. We show that this leads to consistent rankings in all metrics and avoids the penalties of using one saliency map for all metrics. Under this framework, "good" models will perform well in all metrics.

1. Introduction

Humans have a foveated visual system: only a small central part of the retina has high receptor density allowing the perception of the details of a scene. Therefore humans make eye movements to place the high resolution fovea on things they want to see. Understanding where they choose to look is therefore an important component of understanding behaviour.

A long-standing account of bottom-up attentional guidance posits the existence of a "saliency map" (or maps) in the human brain [1], [2]. Here, a saliency map topographically represents importance, usually defined to be local contrast in low-level features such as luminance, color or orientation. Since Itti and Koch formulated this concept into their seminal image-based model [3], the quest of predicting fixations has evolved into a mature field. New models are published on a near-weekly basis with contributions coming mainly from the communities of computer vision and psychology. Over time, the concept of a saliency map has moved away from its origins in low-level feature integration, and now refers more generally to "a map that predicts fixations". In practice, saliency maps are now synonymous with saliency models.

The community uses benchmarks to track progress in the field, of which the MIT Saliency benchmark [4] is the most popular. However, even though there is a clearly accepted benchmark, the community does not agree on which model is the best. Why?

Over the course of decades, a variety of metrics have been proposed to assess the performance of a saliency model. The MIT benchmark evaluates the submission in eight different metrics. Depending on which metric one chooses, the model rankings and performances change dramatically. When formulating a model, the researcher has to decide for which metric the saliency maps should be optimized, knowing that the model will be penalized by other metrics.

A considerable amount of research has analyzed the differences between these metrics, leading to the general conclusion that the metrics measure qualitatively different things [5]–[7]. This is acceptable if one considers the metrics to be different tasks (see below), that may or may not be of interest, but ultimately we might hope that the "best" model would do all tasks well. As it stands however, this will not necessarily be the case for all metrics. A deeper barrier to consistent evaluation is that the metrics use different definitions of how predictions should be formulated in a saliency map. We have recently shown that phrasing saliency models probabilistically and considering the output of a model to be the log probability density removed nearly all disagreement between the metrics [8].

Currently however, a researcher must still decide on a particular saliency map to submit to the benchmark, and therefore will be penalized in other metrics—not because the model is intrinsically bad on those metrics, but because it is being evaluated on a task it was not asked to perform (see Figure 1). For example, AUC expects the model to account for the central viewing bias whereas sAUC will penalize a model that does (and vice versa).

Here, we argue that the fundamental problem underlying these issues is that saliency models and saliency maps are considered to be the same. We propose instead that saliency metrics measure the performance of the model in a spe-

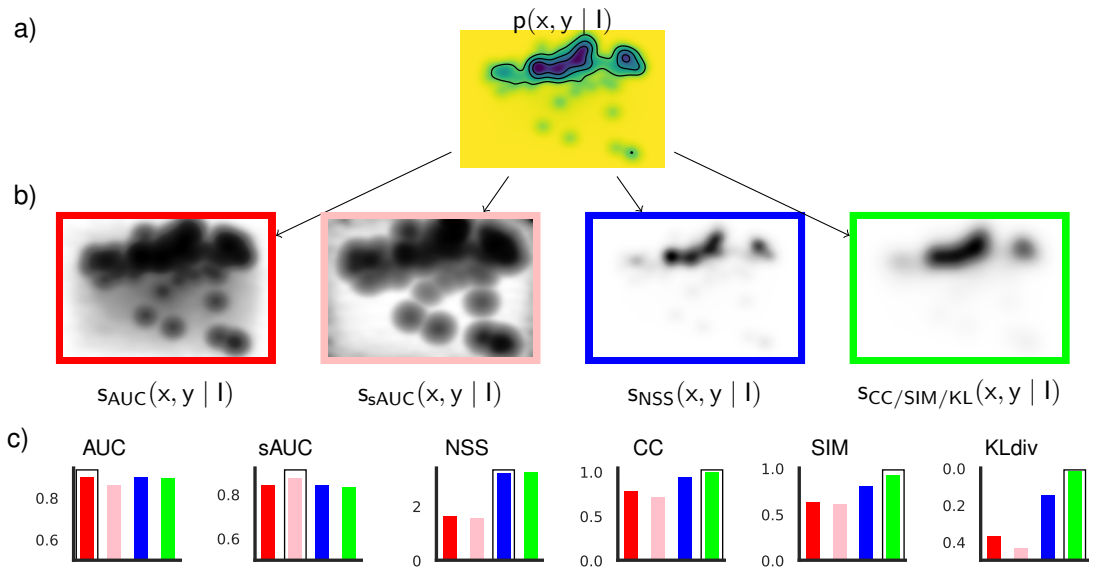

Figure 1: No single saliency map can perform best in all metrics even when the true fixation distribution is known. This problem can be solved by separating saliency models from saliency maps. **a)** fixations are distributed according to a ground truth fixation density $p(x, y \mid I)$ for some stimulus I (see Section 4.5 for details on the visualization). **b)** this ground truth density predicts different saliency maps depending on the intended metric. The saliency maps differ dramatically due to the different properties of the metrics but always reflect the same underlying model. **c)** Performances of the saliency maps from b) under six saliency metrics on a large number of fixations sampled from the model distribution in b). Colors of the bars correspond to the frame colors in c). The predicted saliency map for the specific metric (framed bar) yields best performance in all cases.

cific *task* which saliency models perform by predicting a saliency map that they expect to have the highest rating for this metric. We show that predicting saliency maps given a fixation density is principled for the most influential metrics AUC, sAUC, NSS, CC, SIM, and KL-Div. By considering saliency maps as task-specific predictions, saliency models can avoid the systematic penalties caused by not using the same definition of "saliency map" as expected by the metric (Figure 1).

2. Results

One reason for the metric confusion is that there is no commonly accepted definition of what a saliency model is [8]. For historical reasons, saliency models and saliency maps are usually treated as the same thing (see Introduction). Here we propose and use the following definitions:

1. a *saliency model* predicts a fixation probability density $p(x, y \mid I)$ given an image I.

2. a *saliency metric* measures the performance of a saliency model in a specific task.

3. a *saliency map* $s_{p,\text{task}}(x, y, I)$ is task-specific prediction derived from the model density.

It has been argued before that formulating saliency models as probabilistic models is advantageous (e.g. [8], [9]). In our definition, a saliency model predicts a fixation probability density, that is, the probability $p(x, y \mid I)$ of observing a fixation at a given pixel in a given image [1]. Most model predictions as previously used cannot be interpreted this way.

Note that the three definitions we propose above follow the rational of Bayesian decision theory. The saliency model is a posterior density over all possible events and the saliency metric is a utility function. Based on the posterior density and the utility function, a saliency map is then chosen (or *decided* for) to maximize the expected utility.

2.1. Predicting saliency maps from saliency models

From the predicted fixation density of a model, one can use expected utility maximization to derive the saliency

[1] Note that we use the fixation probability density for single fixations (as in [8]) whereas [9] define a point process density for a whole scanpath.

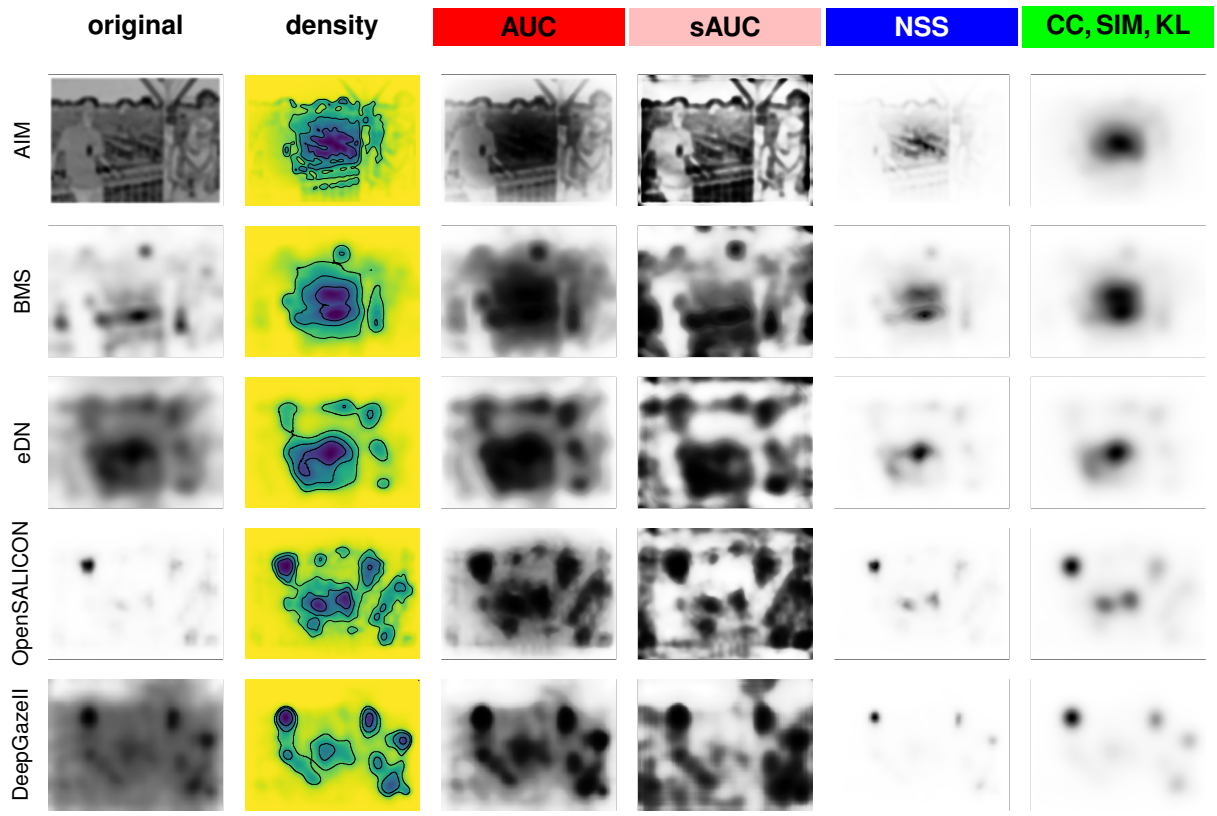

Figure 2: The predicted saliency map for various metrics according to different models, for the same stimulus. For five models (rows) we show their original saliency map (first column), the probability distribution after converting the model into a probabilistic model (second column) and the saliency maps predicted for four different metrics (columns three through six). The predictions of different models for the same metric (column) appear more similar than the predictions of the same model for different metrics (row). In particular, note the inconsistency of the original models (what are typically compared on the benchmark) relative to the per-metric saliency maps. It is therefore difficult to visually compare original model predictions, which have been formulated for different metrics.

map which the model expects to yield highest performance in some metric[2].

Evaluating a saliency metric involves a saliency map $s(x, y \mid I)$ for a stimulus I and ground truth fixation data (x_i, y_i). Therefore, we can phrase a metric as a function $M[s(x, y \mid I); (x_1, y_1), \ldots, (x_n, y_n)]$. Note that some metrics as CC or SIM use an empirical saliency map instead of ground truth fixations. However, the empirical saliency map is always constructed from ground truth fixations, usually by convolving them with a Gaussian. This can be taken to be part of the metric evaluation, as we will demonstrate below. Simplifying notation with $D = (x_1, y_1), \ldots, (x_n, y_n)$, the metric evaluation can be written as

$$M[s(x, y \mid I); D].$$

If we or a model assume that the fixations are distributed according to some distribution $(x_i, y_i) \sim p(x, y \mid I)$ and therefore $D \sim \prod_1^n p(x, y)$, the expected performance of the metric on a saliency map is $\mathbb{E}_D M[s(x, y \mid I); D]$. The model will choose the saliency map which it expects to yield highest performance for the metric M: that is, the solution of

$$\max_{s(x,y|I)} \mathbb{E}_D M[D, s(x, y \mid I)]$$

Solving this optimization problem for a metric M gives us a transformations $p(x, y \mid I) \mapsto s_M(x, y \mid I)$ from fixation densities to predicted metric-specific saliency maps.

[2] Note that the term "metric" is a slight abuse of notation: strictly speaking, a metric measures the distance between two objects and is usually desired to be minimal. However, in saliency, the term "metric" denotes the performance that one wants to maximize (with a few exceptions, e. g., KL-Div and earth mover's distance).

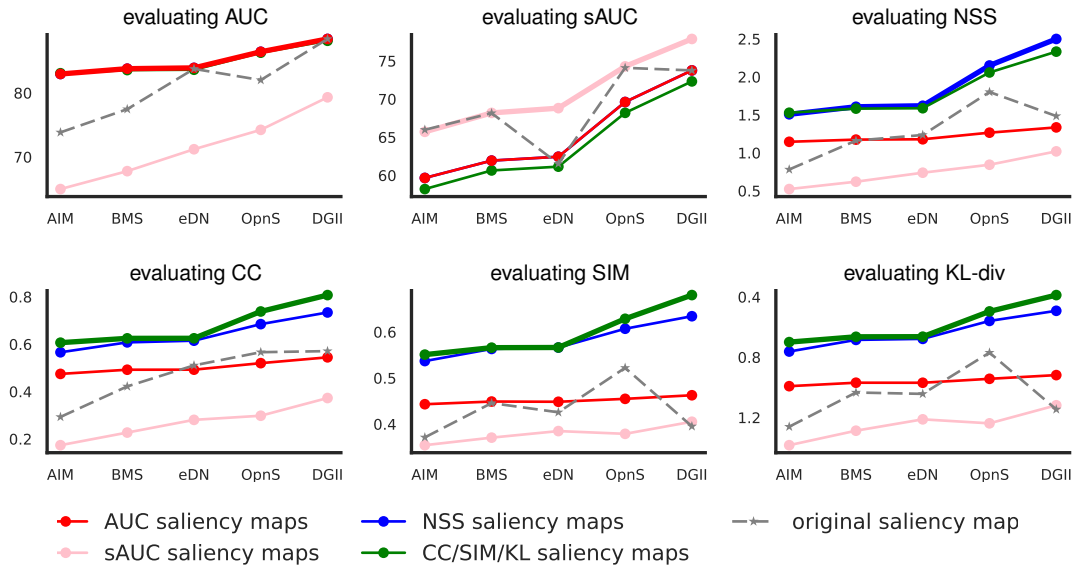

Figure 3: We reformulated several saliency models in terms of fixation densities and evaluated AUC (upper left), sAUC (upper center), NSS (upper right), CC (lower left), SIM (lower center), and KL-div (lower right) on the original saliency maps (dashed line) and the saliency maps predicted by the probabilistic model for the different saliency metrics (solid lines) on the MIT1003 dataset. Saliency maps transformed for a given metric always yield the highest performance (thick line), and for each metric the model ranking is consistent across the different types of saliency maps except the original saliency maps. Note that AUC metrics yield identical results on AUC saliency maps and NSS saliency maps, therefore the blue line is hidden by the red line in the AUC and sAUC plots. OpnS=OpenSALICION, DGII=DeepGaze II.

While this optimization problem might be hard in general, for most commonly-used saliency metrics it can be solved exactly or approximately, as we show below (see supplementary material for further details). Importantly, the methods we outline here are deterministic transformations depending only on the model's density prediction. No optimization using new data is necessary. This can be thought of as the model predicting a saliency map for a given task from its density.

In the following we give exact or approximate solutions for six of the most widely used metrics, including three metrics which operate directly on ground truth fixations (AUC, sAUC, and NSS) and three distribution-based metrics which first convert the ground truth fixations into a empirical saliency map.

AUC, sAUC The AUC-type metrics ("Area Under the Curve") measure the model performance in an 2AFC task where the model has to decide which one of two locations has been fixated. From this observation a very simple calculation shows that the saliency map is the quotient of the fixation distribution by the nonfixation distribution (see supplementary material and [9]). An additional practical consideration is that the MIT benchmark currently only accepts submissions as JPEG images. To compensate for this limited precision and possible JPEG-artefacts, one should additionally histogram-equalize the saliency map.

NSS The Normalized Scanpath Saliency (NSS) performance of a saliency map model is defined to be the average saliency value of fixated pixels in the normalized (zero mean, unit variance) saliency maps. An analytical solution shows that the model should expect the highest NSS score from the predicted fixation density itself (see supplementary material).

CC The correlation coefficient (CC) measures the correlation coefficient between model saliency map and empirical saliency map after normalizing both saliency maps to have zero mean and unit variance. This is equivalent to measuring the euclidean distance between the predicted saliency map and the normalized empirical saliency map. The expected euclidean distance to a random variable is minimized by its expectation value. Therefore the optimal saliency

map with respect to CC is the expected normalized empirical saliency map (see Supplementary material for more details).

Empirical saliency maps are typically computed by blurring observed fixation positions from eye movement data with a Gaussian kernel of a certain size. In this case the expected empirical saliency map is the predicted density convolved with the same Gaussian kernel. The normalization makes the problem analytically intractable. Therefore we use the blured predicted density as an approximation and verify with simulations that the normalization does not affect the expectation value significantly (See supplementary material).

SIM The *Similarity* (SIM) metric normalizes the saliency maps to be probability vectors (nonnegative, unit sum) and sums the pixelwise minimum of two saliency maps. As opposed to the CC-metric, which can be interpreted as measuring the l_2-distance between normalized saliency maps, this effectively measures the l_1-distance between saliency maps. As for CC, this optimization problem is also intractable in general, but this analogy already suggests that the same saliency map as for CC might serve as a good approximation to the optimal saliency map. We verified this with simulations (see supplementary material for more details).

KL-Div The KL-DIV metric computes the Kullback-Leibler divergence between the empirical saliency maps and the model saliency maps by treating both as probability distributions (in the same way as done by SIM). An analytical solution shows that the expected empirical saliency map expects the best performance. As for CC, this is the density blurred by the same kernel size as used for the empirical saliency map.

2.2. Evaluating metric-specific saliency maps

We visualize our approach in Figure 1. In Figure 1a we show a fictional probability distribution for some stimulus. Figure 1b shows the four saliency maps that we predict to be optimal for the six saliency metrics AUC, sAUC, NSS, and CC/SIM/KL-Div. Although they all are predicted by the same model, they appear visually different: while the AUC saliency map is essentially just the normalized density, the sAUC saliency map removes the center bias contribution (see above). The NSS saliency map is exactly the density and shows large areas with very low values. The CC/SIM/KL saliency map, being a blurred version of the density, is much smoother than the NSS saliency map.

In Figure 1c we evaluate the six saliency metrics of interest (six bar plots) on the four predicted saliency maps (colored bars) using fixations sampled from the density in 1a (i.e., the model density is the true density). The rankings of the four saliency maps is highly inconsistent across metrics: even with knowledge of the real fixation distribution, no saliency map can be optimal for all saliency metrics. However, each saliency map is optimal for exactly those metrics for which is has been predicted to be optimal (framed bars). This illustrates our main result: By predicting metric-specific saliency maps in a principled way from fixation densities, one model can perform optimally in all metrics.

Example saliency maps In Figure 2, we show the probability distribution and the predicted saliency maps (columns) for five saliency models (rows) for one example stimulus. Comparing the saliency maps within and between columns, i.e. metrics, one notices that the process of predicting saliency maps for certain metrics has a strong effect on the shape of the saliency maps that is consistent across models. It influences the visual appearance of the saliency map to a larger degree than the actual model does: the AUC and sAUC maps are very high contrast, while the NSS and CC saliency maps have large areas of very little saliency. The CC saliency maps are much smoother than all other saliency maps. It is a quite common technique in the field to compare the saliency maps of different models visually (e.g., see [10], Figure 6; [11], Figure 6; [12], Figure 9). Figure 2 shows that this technique can be very misleading unless the saliency maps are of the same type (e.g. optimized for the same saliency metric).

Comparing model performance In Figure 3 we evaluate the saliency maps of five saliency models (AIM, BMS, eDN, OpenSALICON, DeepGaze II; x-axis) on the same six saliency metrics (subplots). Each line indicates the models' performances in the evaluated metric when using a specific type of saliency map. The dashed lines indicate performance using the models' original saliency maps (i.e. not transformed into true probability densities). The performances are very inconsistent between the different metrics on the original saliency maps. The solid lines indicate the metric performances on the four types of predicted saliency maps (red: AUC, pink: sAUC, blue: NSS, green: CC). In each metric evaluation, the line corresponding to the saliency maps predicted for that specific metric are shown with increased line width. Except for very few borderline cases (see below), the thick lines yield highest performance for all models and metrics, while the other saliency map types often incur severe penalties compared to the thick line. This shows that by predicting metric-specific saliency maps from a predicted density, models can yield high ratings in all metrics.

Interestingly, the AIM model reaches better NSS performance with the CC saliency map than with the NSS saliency

map. This is easy to explain: the AIM model's predicted density improves after blurring. For the better models this effect vanishes. For example, DeepGaze II reaches significantly higher NSS scores with the NSS saliency map than with the CC saliency map and vice versa for the CC metric.

Additionally, we would like to point out that these plots show that as long as one compares all models using the same saliency map type, the metrics agree in their rankings. This generalizes our results from [8], where we have shown this for log-densities.

3. Discussion

Despite much progress in saliency prediction in recent years, comparing saliency models to each other can be confusing due to the large number of benchmarking metrics, giving inconsistent model rankings. Here we argue that benchmarking can be simplified by considering *saliency models* to be probability density predictors, *saliency metrics* to be performance in specific tasks that models perform using that density, and *saliency maps* to be task-specific predictions derived from the model's density. We have shown probabilistic models can predict good saliency maps for the most common saliency metrics: "good models" perform well in many tasks.

Importantly, this task-specific prediction reflects the same underlying model. It is not the case that the model is being re-trained for each task. Rather, the saliency maps we show are derived deterministically from the fixation density predicted by a model. In this way it is possible to fairly compare models by making sure that they are performing the same task, for each task of interest. Conversely, it is misleading to visually compare saliency maps intended for different metrics (see Figure 2)—but this is commonly done in the field ([10]–[12])

Distribution-based metrics like CC, SIM and KL-Div can be harder to use because they depend on how the empirical saliency maps are computed (see supplementary material in [8]). Since we have shown that the optimal saliency maps for these metrics require blurring, the present result shows that these metrics will penalize saliency maps if they correctly predict fixations to be very concentrated. Distribution-based metrics always require the saliency map to be smoother than the true distribution.

Another consequence of the present work is that the eight metrics available on the MIT benchmark can now be seen as a benefit rather than a possible source of confusion. Since each metric assesses a different task, the benchmark would now allow fair comparison over a number of tasks of interest, which may be more or less relevant for certain applications. For example, sAUC is most relevant when one is interested in a model's predictive performance once the center bias is excluded (e.g., in applying to a setting with a different center bias from the MIT1003 training data).

The model comparison via task-specific saliency maps we have outlined here is only possible by phrasing saliency models probabilistically and optimizing them for information [8]. Information gain (equivalently, log-likelihood) is an ideal optimization metric because it reflects all information in the structure of the fixation density, independent of any particular task. Therefore, it should lead to good results in all metrics.

While the saliency maps we have derived give the optimal task-specific saliency map for a given fixation density, it is nevertheless still possible that a given model could do better on a metric with a saliency map not intended for that metric, rather than the task-specific saliency map itself. If the model's density is not the correct one (i.e. does not reflect the data-generating density), then the model can predict suboptimal saliency maps. If the model's density is especially bad, some metrics might even perform better on saliency maps not predicted for this metric than on the one predicted for this metric. For example: if a model's density prediction is too sparse, the AUC metric will perform better on the CC saliency map than it will perform on the actual AUC saliency map. Therefore, actually optimizing model predictions for each specific task may yield insights into the differences between the tasks (by comparing densities optimized for the different metrics). In this case however, we *would* be comparing different models.

Finally, we would like to note that the distinction between saliency models and saliency maps we draw here does not contradict ideas that a "saliency map" or maps may be instantiated in the human brain, as a corollary of bottom-up attentional guidance or an importance map for (e.g.) choosing the next place to fixate in a scene [1], [2], [13]. Our nomenclature is rather independent and intended for saliency model benchmarking.

We will include code for evaluating saliency models as demonstrated in this work in the python library `pysaliency` (available at `https://github.com/matthias-k/pysaliency`).

Conclusion The purpose of this paper was to show how the *same* model can be used to make task-specific predictions, allowing fair comparison between models. In current practice, this would mean submitting different saliency maps for each metric on the MIT benchmark. In the future, we suggest that benchmarks could accept fixation densities (in a high-precision format) and then use these densities to derive model predictions for each metric (saliency maps) according to the transformations we have shown here.

4. Methods

4.1. Image dataset and fixation prediction models

We use the popular benchmarking dataset MIT-1003 ([14]) to compare and evaluate fixation prediction models.

For all evaluated models, the original source code and default parameters have been used unless stated otherwise. The included models are **AIM** [15], Boolean Map-based Saliency **(BMS)** [16], the Ensemble of Deep Networks **(eDN)** [17], **OpenSALICON** [18], and finally **DeepGaze II** [19].

4.2. Phrasing saliency maps probabilistically

To convert an existing saliency-map-based model to a probabilistic model, we used a variant of the method we described in [8]: we fitted a pixelwise monotone nonlinearity and a center bias for each model to yield maximum information gain for the MIT1003 dataset. Unlike in [8] we did not optimize an additional gaussian convolution to smooth the predictions. Since DeepGaze II is already formulated as a probabilistic model, there was no need to convert this model. For showing the "original saliency map" we used the log density in this case.

4.3. Saliency metrics

AUC: We use all pixels as nonfixations. As thresholds we use the combined saliency values of all fixations and nonfixations. **sAUC**: We use the fixations of all other images as nonfixations. As for AUC, we use the combined saliency values of all fixations and nonfixations as thresholds. **NSS** computes the mean saliency of fixation locations after normalizing the saliency map to have zero mean and unit variance. **CC**: As suggested for this dataset ([14]), we convolve the fixation maps of the ground truth fixations with a gaussian kernel with $\sigma = 35px$ to compute empirical saliency maps. **SIM**: We use the same empirical saliency maps as for CC. To convert a saliency map to a probability distribution, we check whether any values of the saliency map are negative. If so, we substract the minimal value from the saliency to make it non-negative. Afterwarts we divide the saliency by the sum of all values. **KL-Div**: We use the same empirical saliency maps as for CC and the same normalization procedure as for SIM.

4.4. Predicting saliency maps from densities

From a probability density over an image we derive four types of saliency maps:

AUC saliency maps are created by equalizing the probability density to yield a uniform histogram over all pixels.

sAUC saliency maps are created by dividing the probability density by the center bias density and again equalizing the saliency map to yield a uniform histogram over all pixels. The center bias density was estimated using a Gaussian kernel density estimate over all fixations from the MIT1003 dataset and crossvalidated across images.

NSS saliency maps are simply the probability density

CC/SIM/KL saliency maps are calculated by convolving the probability density with a gaussian kernel with $\sigma = 35px$.

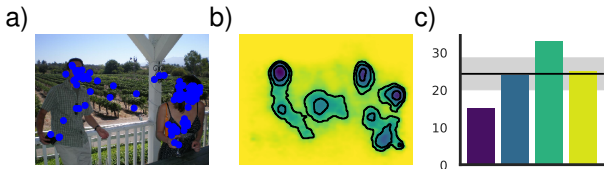

Figure 4: Visualizing fixation densities: **a)** an example stimulus with $N = 97$ ground truth fixations. **b)** DeepGaze II predicts a fixation density for this stimulus. The contour lines separate the image into four areas of decreasing probability density such that each area has the same total probability mass. **c)** The number of ground truth fixations in each of the four areas. The model expects the same number of fixations for each area (horizontal line: 24.25 fixations for N fixations total). The gray area shows the expected standard deviation from this number. DeepGaze II overestimates the how peaked the density is: there are too few fixations in darkest area. Vice versa, it misses some probability mass in the second to last area. However, the large error margin (gray area) indicates that substantial deviations from the expected number of fixations are to be expected.

4.5. Visualizing probability densities

Visualizing two dimensional densities is harder than it appears to be at the first glance: Although the absolute density values have a very precise meaning, it is hard to read substantially more than the ranking of the values and maybe a very rough idea about the peakyness of the distribution from a color map. When visualizing two dimensional probability densities, we add three contour lines separating the image into four areas of decreasing probability density such that each area has the same total probability mass (i.e. the density predicts each area to receive the same number of fixations, see Figure 4b). If the darkest area is very small, this means the density predicts on fourth of the fixations to be clustered in a very small area. If all areas are roughly of the same size, the density is nearly uniform. Comparing the number of fixations in each area can serve as a simple heuristic to asses a model's quality (see Figure 4c).

References

[1] A. M. Treisman and G. Gelade, "A feature-integration theory of attention," *Cognitive psychology*, vol. 12, no. 1, pp. 97–136, 1980.

[2] C. Koch and S. Ullman, "Shifts in selective visual attention: Towards the underlying neural circuitry," *Human Neurobiology*, vol. 4, pp. 219–227, 1985. [Online]. Available: https://cseweb.ucsd.edu/classes/fa09/cse258a/papers/koch-ullman-1985.pdf.

[3] L. Itti, C. Koch, and E. Niebur, "A model of saliency-based visual attention for rapid scene analysis," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 20, no. 11, pp. 1254–1259, 1998. DOI: 10.1109/34.730558.

[4] Z. Bylinskii, T. Judd, F. Durand, A. Oliva, and A. Torralba, *Mit saliency benchmark*, http://saliency.mit.edu/.

[5] N. Wilming, T. Betz, T. C. Kietzmann, and P. König, "Measures and limits of models of fixation selection," *PLOS ONE*, vol. 6, no. 9, e24038, Dec. 9, 2011, ISSN: 1932-6203. DOI: 10.1371/journal.pone.0024038. [Online]. Available: http://journals.plos.org/plosone/article?id=10.1371/journal.pone.0024038.

[6] N. Riche, M. Duvinage, M. Mancas, B. Gosselin, and T. Dutoit, "Saliency and human fixations: State-of-the-art and study of comparison metrics," IEEE, Dec. 2013, pp. 1153–1160, ISBN: 978-1-4799-2840-8. DOI: 10.1109/ICCV.2013.147. [Online]. Available: http://ieeexplore.ieee.org/document/6751253/.

[7] Z. Bylinskii, T. Judd, A. Oliva, A. Torralba, and F. Durand, "What do different evaluation metrics tell us about saliency models?" *ARXIV:1604.03605 [cs]*, Apr. 12, 2016. arXiv: 1604.03605. [Online]. Available: http://arxiv.org/abs/1604.03605.

[8] M. Kümmerer, T. S. A. Wallis, and M. Bethge, "Information-theoretic model comparison unifies saliency metrics," *Proceedings of the National Academy of Sciences*, vol. 112, no. 52, pp. 16 054–16 059, Dec. 2015. DOI: 10.1073/pnas.1510393112. [Online]. Available: http://dx.doi.org/10.1073/pnas.1510393112.

[9] S. Barthelmé, H. Trukenbrod, R. Engbert, and F. Wichmann, "Modelling fixation locations using spatial point processes," *Journal of Vision*, vol. 13, no. 12, 2013. DOI: 10.1167/13.12.1.

[10] M. Cornia, L. Baraldi, G. Serra, and R. Cucchiara, "Predicting human eye fixations via an LSTM-based saliency attentive model," *ARXIV:1611.09571 [cs]*, Nov. 29, 2016. arXiv: 1611.09571. [Online]. Available: http://arxiv.org/abs/1611.09571.

[11] A. Borji, D. N. Sihite, and L. Itti, "Quantitative analysis of human-model agreement in visual saliency modeling: A comparative study," *IEEE Transactions on Image Processing*, vol. 22, no. 1, pp. 55–69, Jan. 2013, ISSN: 1057-7149, 1941-0042. DOI: 10.1109/TIP.2012.2210727. [Online]. Available: http://ieeexplore.ieee.org/document/6253254/.

[12] A. Borji and L. Itti, "State-of-the-art in visual attention modeling," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 1, pp. 185–207, Jan. 2013, ISSN: 0162-8828. DOI: 10.1109/TPAMI.2012.89.

[13] Z. Li, "A saliency map in primary visual cortex," *Trends in cognitive sciences*, vol. 6, no. 1, pp. 9–16, 2002.

[14] T. Judd, K. Ehinger, F. Durand, and A. Torralba, "Learning to predict where humans look," in *Computer Vision, 2009 IEEE 12th international conference on*, IEEE, 2009, pp. 2106–2113.

[15] N. D. Bruce and J. K. Tsotsos, "Saliency, attention, and visual search: An information theoretic approach," *Journal of vision*, vol. 9, no. 3, 2009.

[16] J. Zhang and S. Sclaroff, "Saliency detection: A boolean map approach," in *Computer Vision (ICCV), 2013 IEEE International Conference on*, IEEE, 2013, pp. 153–160.

[17] E. Vig, M. Dorr, and D. Cox, "Large-scale optimization of hierarchical features for saliency prediction in natural images," in *Computer Vision and Pattern Recognition, 2014. CVPR'14. IEEE Conference on*, IEEE, 2014.

[18] C. Thomas, "Opensalicon: An open source implementation of the salicon saliency model," *CoRR*, vol. abs/1606.00110, 2016. [Online]. Available: http://arxiv.org/abs/1606.00110.

[19] M. Kümmerer, T. S. A. Wallis, and M. Bethge, "DeepGaze II: Reading fixations from deep features trained on object recognition," *ARXIV:1610.01563 [cs, q-bio, stat]*, Oct. 5, 2016. arXiv: 1610.01563. [Online]. Available: http://arxiv.org/abs/1610.01563.

5. Supplementary Material

5.1. Solving the optimization problems

AUC, sAUC: The AUC-type metrics measure the model performance in an 2AFC task where the model has to decide which one of two locations has been fixated (See Figure 5 for a visual proof). Let's call the model's fixation distribution $p_{\text{fix}}(x, y)$, the nonfixation distribution

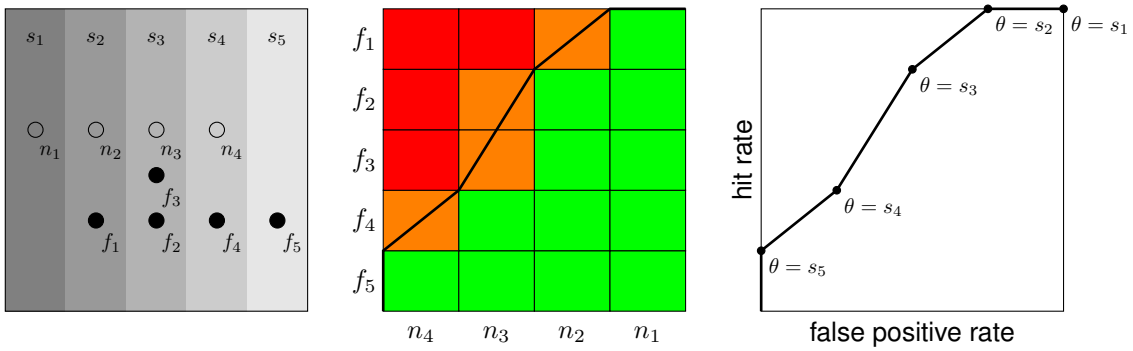

Figure 5: AUC metrics measure the performance of the saliency map in a 2AFC task where the saliency values of two locations are used to decide which of these two locations is a fixation and which is a nonfixation. **a)** An example saliency map is shown consisting of five saliency values ($s_1 < \cdots < s_5$) and with five fixations ($f_1, \ldots, f_5$) and four nonfixations ($n_1, \ldots, n_4$). **b)** The performance in the 2AFC task can be calculated by going through all fixation-nonfixation pairs (f_i, n_j): The saliency map decides correct if the saliency value of f_i is greater than n_j (green), incorrect if it is smaller (red) and has chance performance if the values are equal (orange). Below the thick line are all correct predictions (green) and half of the chance cases (orange). **c)** The ROC curve of the saliency map with respect to the given fixations and nonfixations. For each threshold θ all values of saliency value greater or equal to θ are classified as fixations. Comparing b) and c) shows that the area under the curve in c) is exactly the performance in the 2AFC task in b).

$p_{\text{nonfix}}(x, y)$ (which is uniform for AUC and the image independent center bias for sAUC) and denote the two locations by (x_1, y_1) resp. (x_2, y_2). The 2AFC task boils down to deciding whether these points are sampled from $p_{\text{fix}} \times p_{\text{nonfix}}$ for from $p_{\text{nonfix}} \times p_{\text{fix}}$. The likelihoods of the two points given these two distributions are $p_{\text{fix}}(x_1, y_1)p_{\text{nonfix}}(x_2, y_2)$ resp. $p_{\text{nonfix}}(x_1, y_1)p_{\text{fix}}(x_2, y_2)$. The model expects optimal performance by choosing the distibution which has higher likelihood, or equivalently, the point for which $p_{\text{fix}}(x, y)/p_{\text{nonfix}}(x, y)$ has the higher value. Therefore the model should expect the saliency map $p_{\text{fix}}(x, y)/p_{\text{nonfix}}(x, y)$ to yield highest performance. In the special case of the standard AUC metric, p_{n} is constant and the saliency map boils down to p_{f}.

There is one additional technicality worth considering: the AUC metrics don't care about absolute saliency values or even about the value of the difference between saliency values. They only care about the ranking of the saliency values, or equivalently: about the sign of the difference between saliency values. Therefore, if the saliency maps are stored with reduced precision (e.g. the MIT saliency benchmark expects submissions to be 8bit JPEG files), the saliency map should be histogram-equalized to make sure the sign of saliency value differences stays constant as often as possible.

NSS: The *normalized scanpath saliency* (NSS) of a saliency map model is defined to be the average saliency value of fixated pixels in the normalized (zero mean, unit variance) saliency maps. The best possible saliency map with respect to the NSS metric can be deduced from a probabilistic formulation of fixations: Given an image with N pixels let the probability for a single fixation falling onto pixel i be p_i. Then the expected NSS of a saliency map $q = (q_1, \ldots, q_N)$ with $\frac{1}{N}\sum_i q_i = \bar{q} = 0$, $\sum_i^N q_i^2 = \|q\|_2^2 = 1$ is

$$\mathbb{E}NSS(q) = \sum_i^N p_i \cdot q_i = \langle p, q \rangle.$$

Finding the saliency map with the best possible NSS is equivalent to finding the solution of the problem

$$\max \langle p, q \rangle \quad s.t.\ \bar{q} = 0, \|q\|^2 = 1$$

Since $q \mapsto q' = \bar{p} + \alpha q$ with $\alpha = \sqrt{\|p\|^2 - 1/N}$ induces a maximum-preserving bijection between $\{q \mid \bar{q} = 0, \|q\|^2 = 1\}$ and $\{q' \mid \bar{q}' = \bar{p} = 1/N, \|q\|^2 = \|p\|^2\}$, we can look for a solution to

$$\max \langle p, q' \rangle \quad s.t.\ \bar{q}' = \bar{p}, \|q'\|^2 = \|p\|^2$$

instead (and normalize q afterwarts, if we want to have the normalized saliency map). Because of $\langle x, y \rangle = \frac{1}{2}(\|x\|^2 + \|y\|^2 | - \|x - y\|^2)$, the maximum under these conditions is at the same time the minimum of $\|p - q\|^2$, which is obviously p.

Therefore, the best possible saliency map with respect to NSS is the density of the fixation distribution.

CC: The CC metric is slighty more involved since it does not operate directly on fixations but on the empirical saliency maps derived from the fixations. CC measures the

correlation coefficient between the model saliency map and the empirical saliency map after both have been normalized to zero mean and unit variance.

Because of $\langle x, y\rangle = \frac{1}{2}(\|x\|^2 + \|y\|^2 - \|x-y\|^2)$, under the normalization conditions maximizing the correlation coefficient is equivalent to minimizing the euclidean distance. Since the euclidean distance to a random variable is minimized by its expectation value, this shows that the optimal point estimate for the CC metric is the expected normalized empirical saliency map.

Therefore, predicting the optimal saliency map for CC crucially depends on how the empirical saliency maps are computed. Usually, they are computed by putting a Gaussian of size σ at each fixation location or equivalently, convolving the fixation maps by a Gaussian convolution with size σ. In this case the expected empirical saliency map would be

$$\begin{aligned}\mathbb{E}_{x_i\sim p}\frac{1}{N}\sum_i^N G_\sigma(x) &= \frac{1}{N}\sum_i^N \mathbb{E}_x \sim pG_\sigma(x)\\ &= \frac{1}{N}\sum_i^N G_\sigma * p\\ &= G_\sigma * p,\end{aligned}$$

that is, the density blurred with a Gaussian kernel of size σ.

Unfortunately, we don't need the expected empirical saliency map but the expected normalized saliency map (which is not the same as the normalized expected saliency map). The normalization involves subtracting the mean and dividing by the standard deviation. Subtracting the mean is irrelevant, however, normalizing the variance is nonlinear. Nevertheless, normalizing by the variance effectively just changes the weight by which the different empirical saliency maps are averaged in the expectation value. As long as the variances of the different empirical saliency maps don't differ too much, this won't have much of an effect and our simulations suggest that this is the case (Figure 6). Therefore, as an approximation to the expected normalized empirical saliency map, we use the expected saliency map in this paper, which is computed by convolving the expected density by a Gaussian.

Obviously, if more involved techniques are used to compute the empirical saliency maps (e.g. cross validation of the kernel size as in [8]), then the expected empirical saliency map is harder or impossible to calculate analytically. However, one can still approximate it numerically by sampling empirical saliency maps from the expected fixation distribution and averaging them.

SIM: The *similarity-metric* normalizes the saliency map and the empirical saliency map to be a probability distribution (i.e. summing up to one) and then sums over the pixelwise minimum of the two saliency maps. Since

$$\begin{aligned}\sum_i \min(p_i, q_i) &= \sum_i \frac{1}{2}\left(p_i + q_i - |p_i - q_i|\right)\\ &= 1 - \frac{1}{2}\|p - q\|_1,\end{aligned}$$

the similarity is essentially just measuring the l_1-distance between the two saliency maps (where $p = (p_1, \ldots, p_N)$ and q indicate saliency maps which are flattened into vectors to ease notation). This makes SIM a close relative of the CC-metric which is measuring the l_2-distance as shown above). Let $p = (p_1, \ldots, p_N)$ with $p \geq 0$, $\sum_i p_i = 1$ denote the random variable which represents the empirical saliency map and q with $q \geq 0$, $\sum_i q_i = 1$ the model saliency map (for simplicity we are flattening the two dimensional saliency maps). Then we are looking for the q which minimizes $\mathbb{E}_p\|q-p\|_1$ under the constraint that $q \geq 0$ and $\|q\|_1 = 1$.

Ignoring the constraint, the l_1-distance is minimized by the median while the l_2-distance is minimized by the mean. If a distribution is symmetric, the mean and the median are the same and we might hope this property holds at least approximately for the values of the empirical saliency maps, especially given that they are additionally constrained to be non-negative and have unit sum. We checked this in simulations and found that while the difference is larger than for CC, it still is not very large.

KL-DIV: The KL-DIV metric computes the Kullback-Leibler divergence between the empirical saliency maps and the model saliency maps by treating both as probability distributions (in the same way as done by SIM). Therefore, unlike for most other metrics, in KL-DIV lower values are better. Let $p = (p_1, \ldots, p_N)$ with $p \geq 0$, $\sum_i p_i = 1$ denote the random variable which represents the empirical saliency map and q with $q \geq 0$, $\sum_i q_i = 1$ the model saliency map (for simplicity we are flattening the two dimensional saliency maps). Then we are looking for the q which minimizes $\mathbb{E}_p KL[p, q]$. Since

$$\begin{aligned}\mathbb{E}_p[KL[p, q]] &= \mathbb{E}_p\left[\sum_i p_i \frac{\log p_i}{\log q_i}\right]\\ &= \mathbb{E}_p\left[\sum_i p_i \log p_i\right] - \sum_i \mathbb{E}_p[p_i]\log q_i,\end{aligned}$$

this is equivalent to finding the maximum of $\sum_i \mathbb{E}_p[p_i]\log q_i$, which is again equivalent to finding the minimum of $\sum_i \mathbb{E}_p[p_i]\log \mathbb{E}_p[p_i] - \sum_i \mathbb{E}_p[p_i]\log q_i = KL[\mathbb{E}_p[p], q]$. This is obviously minimized by $q = \mathbb{E}_p[p]$, the expected empirical saliency map.

a) b) c)

d)

kernel size / fixations	CC scores (mean empirical/mean normalized empirical)				
	1	10	20	30	
1	0.0127/0.0127	0.0632/0.0633	0.1210/0.1214	0.2227/0.2242	0.3066/0.3100
10	0.0412/0.0412	0.1928/0.1929	0.3488/0.3491	0.5630/0.5641	0.6855/0.6874
50	0.0920/0.0920	0.3960/0.3960	0.6120/0.6121	0.7728/0.7731	0.8198/0.8206
100	0.1295/0.1295	0.5191/0.5191	0.7307/0.7307	0.8464/0.8465	0.8707/0.8709
200	0.1817/0.1817	0.6509/0.6509	0.8313/0.8314	0.9074/0.9074	0.9203/0.9204

Figure 6: Predicting optimal saliency maps for the CC metric: Starting from a density (a) we sampled 100000 sets of either 1, 10 or 100 fixations and used them to create empirical saliency maps. Using these empirical saliency maps, we calculated the mean empirical saliency map (shown for 10 fixations per empirical saliency map in (b)). Additionally, we normalized the empirical saliency maps to have zero mean and unit variance to compute the mean normalized empirical saliency map (c) which is optimal with respect to the CC metric. Then we sampled another 100000 empirical saliency maps from the original density and evaluated CC scores of the mean empirical and mean normalized empirical saliency maps (d). The mean normalized saliency map yields slighly higher scores in all cases but the difference to the mean empirical saliency map is tiny, indicating that the expected empirical saliency map is a very good approximation of the optimal saliency map for the CC metric.

a) b) c)

d)

fixations / sample	SIM scores	
	mean empirical	optimal
1	0.0815	0.0821
10	0.3657	0.3705
100	0.7499	0.7530

Figure 7: Predicting optimal saliency maps for the SIM metric: Starting from a density (a) we sampled 20000 sets of either 1, 10 or 100 fixations and used them to create empirical saliency maps. Using these 20000 empirical saliency maps, we calculated the mean empirical saliency map (shown for 10 fixations per empirical saliency map in (b)). Additionally, we estimated the optimal saliency map for the SIM metric by using projected gradient descend to find the saliency map which yields highest SIM score over all these empirical saliency map (shown for 10 fixations per empirical saliency map in (c)). Then we sampled another 100000 empirical saliency maps from the original density and evaluated SIM scores of the mean expected and optimal saliency maps (d). The optimal saliency map yields slighly higher scores in all cases but the difference to the mean empirical saliency map is not large, indicating that the expected empirical saliency map is a good approximation of the optimal saliency map for the SIM metric.